

Experimental Longtermist Jurisprudence

Eric Martínez and Christoph Winter

Recent scholarship has revealed a seemingly stark mismatch between the value of future generations and the lack of protection afforded to them under present legal systems. Although climate change, pandemics, nuclear war, and artificial intelligence impose greater threats to the future of humanity than any previous risk (Ord, 2020), legal systems fail to grant future generations democratic representation in the legislature, standing to bring forth a lawsuit in the judiciary, and serious consideration in cost-benefit analyses in the executive. What is the source of this disconnect, is it justified, and—to the extent that it is not justified—what might one do about it?

Here we discuss how a new research field within experimental jurisprudence—which we refer to as *experimental longtermist jurisprudence*—might help address these questions and in turn help determine the appropriate level and form of legal protection to future generations.

The chapter is divided into three sections. In Section 10.1, we provide an overview of the substantive and methodological underpinnings of experimental longtermist jurisprudence. In Section 10.2, we introduce three research programs within experimental longtermist jurisprudence, and in Section 10.3, we discuss the normative implications of each of these research programs.

10.1 Foundations of Experimental Longtermist Jurisprudence

Experimental longtermist jurisprudence (XLJ) is, in essence, a form of experimental jurisprudence¹ concerned with issues relating to the long-term future. The main substantive and methodological underpinnings of experimental longtermist jurisprudence are

1. the view that one should be particularly concerned with ensuring that the long-run future goes well (*longtermism*);
2. the view that said concern ought to extend to the legal system (*legal longtermism*); and
3. the practice of using experimental methods as a means of evaluating the validity and implications of legal longtermism (*experimental legal longtermism*).

Here, we discuss each of these foundations in turn.

10.1.1 The Philosophical Foundations of Longtermism

The set of philosophical theories associated with the view that one should be particularly concerned with ensuring that the long-run future goes well has been referred to as longtermism (MacAskill, 2022).² Longtermism is based on three main assumptions—one normative, two empirical—which we briefly detail here in turn.³

The first assumption, often referred to as the normative assumption, is that when assessing the moral value of our actions, all consequences matter equally—independent of when, where, or how they occur.⁴ For example, just as many of the most influential thought experiments in contemporary moral philosophy have argued or implied that the welfare of someone living far away geographically ought to be valued the same as the welfare of someone living close by (e.g., Singer, 1972; Rawls, 1971; Pogge, 1989; Beitz, 1983), according to longtermism, so too should the welfare of someone living far away in the future be valued the same as the welfare of someone living right now. Additionally, longtermism also argues (in a similar vein as other mainstream philosophical views (see, e.g., Pettit & Smith, 2000; cf. Parfit, 1984, pp. 24–5)) that, *ceteris paribus*, indirect consequences ought to be given the same value as direct consequences, and unintended consequences ought to be given the same value as intended consequences (Greaves & MacAskill, 2019, p. 6). For example, a direct consequence of distributing insecticide-treated bed nets in sub-Saharan Africa is a reduction of malaria incidents and child mortality (Pryce, Richardson, & Lengeler, 2018). However, the fact that these consequences are direct does not on its own imply that these consequences “matter” more *per se* than some of the indirect consequences of distributing insecticide-treated bed nets, such as improved education (Kuecken, Thuilliez, & Valfort, 2014) and increased GDP growth (Gallup & Sachs, 2001; Sachs & Malaney, 2002). Nor does it matter whether said GDP growth or improved education was less intended relative to reducing malaria incidents and child mortality. If this is right, such that all consequences matter equally across time and space, then insofar as future generations exist, it follows that future generations are of equal value in principle as the current generation.⁵

The second assumption states that, *in expectation*, the future is vast in size—that is, it is likely to consist of at least a fairly large number of future generations, each consisting, on average, of a fairly large number of individuals, such that the number of individuals⁶ living in the future will collectively be far greater than the number of individuals living in the present. One way to estimate the future lifespan of humans is by extrapolating from the typical lifespan of a mammalian species, estimated to be anywhere from 600,000 years (Barnosky et al., 2011; Ceballos, 2015) to 1.7 million years (Foote & Raup, 1996). Since *Homo sapiens* is estimated to be 300,000 years old (see Galway-Witham & Stringer, 2018; Schlegel et al., 2017), this would suggest that *Homo sapiens* has 300,000 to 1.4 million years of potential flourishing before extinction. Given that *Homo sapiens* in many ways might be considered more successful than other mammalian species and, in particular, appears less vulnerable to the typical threats of extinction as other mammals,⁷ some experts have pointed out that this estimate may be overly conservative (Greaves & MacAskill, 2019). If so, humans might instead be expected to survive for as long as the earth remains habitable (anywhere between

.9 and 1.5 billion years: Caldeira & Kasting, 1992), even setting aside the prospect of leaving earth and colonizing other habitable systems (Beckstead, 2014), in which case the upper bound, however unlikely, would be as high as quintillions of years (the estimated expected end of the universe: Adams & Laughlin, 1997).

In addition to the expected number of future generations, it also stands to reason that for most (if not all) future generations, there will, in expectation, be greater numbers of people living at any given time than there are presently. After all, the number of people who are living now is estimated to be more than ten times higher than the number of people who were living 200 years ago (7 billion versus 600 million, respectively: Roser, Ritchie, & Ortiz-Ospina, 2019), and current projections estimate that future generations will likewise be greater, even if one assumes that global population growth will slow at a certain point in the near future (United Nations, 2019). In combination with the first assumption (i.e., that consequences affecting each of these individuals would, all else equal, matter just as much as those affecting individuals living in the present), this would imply that insofar as we *can* positively influence the experiences of future generations, their sheer size and value dictates that we *ought* to protect them.

The third assumption states that there are feasible and predictable ways to positively influence the experiences of future generations. For example, while this assumption may initially seem less plausible than the previous two assumptions, given the apparent impracticality of influencing the future in ways that are reasonably foreseeable, longtermists have pointed to examples of both (a) historical trends that have had long-lasting effects on the trajectory of human civilization (e.g., religious values and the implementation of certain legal systems)⁸ and (b) predictable and feasible ways of influencing the future, particularly with respect to existential risks associated with advanced artificial intelligence, extreme climate change, and synthetic biology (Winter et al., 2021; see also Section 10.2.1 *infra*). Together with the previous assumptions, this implies that not only should we value and protect future generations in *principle*, but that we can—and therefore should—protect their interests in *practice*.⁹

10.1.2 Legal Longtermism

One of the primary means through which we might conceivably protect future generations—and the one with which XLJ is chiefly concerned—is via the legal system. The set of views associated with the claim that law and legal institutions ought to protect those in the far future can be referred to as *legal longtermism*.¹⁰ Normatively, the premises associated with legal longtermism are similar, if not identical, to those of philosophical longtermism (depending on the extent to which one believes that those with moral value warrant legal consideration). Empirically, the assumptions are also similar, with the additional supposition that there are feasible and predictable *legal* mechanisms to protect future generations.

That said, given that many of the examples cited in support of the feasibility assumption of philosophical longtermism relate to the legal system, as a practical matter one who accepts the assumptions of philosophical longtermism may automatically accept the premises of legal longtermism. In particular, some of the long-lasting, intergenerational effects of legal systems cited to by legal longtermists include

(a) the medieval establishment of the common law's continued influence on the laws governing Great Britain and its former colonies (Berman, 1985), (b) the long-lasting influence of Roman law on many civil-code systems (Watson, 1991), and (c) and the persistence of Eastern legal institutions (Kuran, 2011; Cheng, Rosett, & Woo, 2003). Legal longtermists have further argued that legal interventions could play a significant role in mitigating some of the catastrophic risks highlighted earlier and ensuring a more positive long-term trajectory (Winter et al., 2021).

As it currently stands, however, legal systems provide hardly any legal protection to future generations. Legal institutions have been and continue to be very short-term oriented, with policy making geared toward solving contemporary issues and with democratic, legislative, and judicial processes generally reserved exclusively for the current generation (cf. ACE Project, 2021; John & MacAskill, 2020). While a few attempts have been made to provide legal protection to future generations, even fewer legal mechanisms have been successfully implemented as a result, and those implemented so far have been largely ineffectual (see, e.g., Araújo & Koessler, 2021).

The primary substantive aims of XLJ relate largely to (a) understanding the source of this disconnect, such as by studying the nature of people's beliefs regarding legal longtermism, how the theory of longtermism coheres with people's ordinary concept of rights and duties, and what contributes to those beliefs and (b) deriving normative implications based on that understanding. The methodological framework for satisfying these aims is detailed in the next section.

10.1.3 Experimental Legal Longtermism

Like other forms of experimental jurisprudence, XLJ employs methods traditionally associated with the field of experimental psychology to explore substantive questions traditionally associated with the field of jurisprudence.¹¹ XLJ can be thought of as an experimental branch of longtermist jurisprudence, with the goals of (a) uncovering the cognitive underpinnings of beliefs relevant to the jurisprudential framework laid out in the previous section and (b) advancing legal, philosophical, and policy arguments on the basis of those findings. In the case of XLJ, these methods include both surveys and controlled experiments. Here, we briefly discuss each of these in turn.

Administering surveys in the context of XLJ research, as in other, more traditional forms of empirical legal studies research, involves straightforwardly asking people questions regarding the aspects of legal longtermism or some legal longtermism-relevant issue. Such research may be aimed at either testing a specific hypothesis or merely gathering information. For example, a survey might evaluate the level of acceptance of different aspects of legal longtermism by identifying which arguments in favor of and against legal longtermism resonate with participants and for what reasons. The results may fuel other hypotheses and be used as inputs into other XLJ research.

Controlled experiments, on the other hand, involve indirectly examining people's views regarding legal longtermism by, for example, asking participants questions about carefully controlled stimuli. Such stimuli are often in the form of contrastive vignettes, in which particular aspects of a situation are systematically manipulated to identify the psychological processes underlying certain concepts, intuitions, or judgments relevant

to legal longtermism. In contrast to the survey method, the primary descriptive aim of the controlled-experiment technique is to better understand the cause of people’s beliefs as opposed to the beliefs themselves.

Unlike conventional surveys and controlled experiments, however, XLJ is concerned not only with descriptive questions of fact but also with normative questions of legal philosophy, doctrine, and policy. Thus, XLJ can be considered a two-step process. In Step 1, a researcher takes the role of a cognitive scientist, trying to gain insight into a legal longtermist-relevant feature of the human mind via an experimental study. In Step 2, a researcher takes the role of a philosopher, legal theorist, lawyer, and/or policy maker, reasoning about the normative implications of the experimental findings uncovered at Step 1.

XLJ can also be conceptualized as containing three separate but interdependent levels of abstraction, corresponding to the three possible sets of normative implications alluded to in the previous paragraph:

1. The *philosophical level*, concerned with determining whether and to what extent future generations ought to be provided legal protection according to an ideal legal system;
2. The *doctrinal level*, concerned with determining to what extent and how future generations ought to be provided legal protection according to the doctrines of the current legal system;
3. The *applied level (or policy level)*, concerned with determining which legal mechanisms and instruments ought to be prioritized and/or implemented so as to provide the appropriate level of legal protection to future generations.

In the remainder of the chapter, we first provide an overview of Step 1 of XLJ research at each of the three levels (Section 10.2) and then turn to an overview of Step 2 of XLJ research at each of the three levels (Section 10.3). For a concise overview of the different levels of abstraction at each of the two steps, see Table 10.1.

Table 10.1 Experimental Longtermist Jurisprudence at Three Levels of Abstraction. © Eric Martínez and Christoph Winter

	Level I: Philosophical Level	Level II: Doctrinal Level	Level III: Applied Level
Step 1: Descriptive Aims	What are people’s general beliefs about the concept of legal longtermism, and why people hold those beliefs?	To what extent does longtermism map onto people’s understanding of law, the legal system, and legally relevant concepts?	What are people’s intuitions about how more specific areas of law and concrete legal mechanisms could address more concrete long-term challenges?
Step 2: Normative Aims	How and to what extent should future generations be provided legal protection (independent of existing legal doctrine)?	How and to what extent should future generations be provided legal protection (according to existing legal doctrine)?	Which legal mechanisms and instruments can and ought to be implemented so as to provide the appropriate level of legal protection to future generations?

10.2 Experimental Longtermist Jurisprudence at Three Levels of Abstraction

As discussed earlier, XLJ can be conceptualized as containing three interrelated levels of abstraction, including what we refer to as the philosophical level, the doctrinal level, and the applied level. Here we discuss each of these levels of abstraction in turn with respect to Step 1 of XLJ research (empirical aims and methodology), while in the final third of the chapter we focus on Step 2 of XLJ research (normative implications).

10.2.1 Philosophical-Level XLJ

Step 1 of XLJ research at the philosophical level involves investigating (a) people's general beliefs about legal longtermism (i.e., the claim that we ought to provide legal protection to those in the far future, as well as the three underlying premises of that claim) and (b) why people¹² hold those beliefs. Here we discuss each of these two types of philosophical-level XLJ methods in turn, in terms of both existing literature and future directions.

With regard to (a) beliefs about legal longtermism, for example, Martínez and Winter (2021a) surveyed a set of over 500 legal academics from around the English-speaking world regarding their views on the desirability and feasibility of using the legal system to protect future generations and influence the long-term future. In terms of desirability, legal academics rated their desired level of legal protection for future generations as several times higher than their perceived current level of legal protection afforded to future generations, and roughly equal to the current level of legal protection afforded to humans living in the present. Moreover, the difference between the desired and current level of legal protection was rated as higher for future generations than for other neglected groups, such as non-human animals. Martínez and Winter (2021b) found similar results in a set of over 1000 US American adults, while Martínez and Winter (in press a) found similar results in a cross-cultural survey of approximately 3000 adults in Australia, Canada, Chile, Japan, Mexico, Spain, South Africa, South Korea, United Kingdom, and the United States (see also Martínez and Winter, in press b).

In terms of feasibility, the vast majority of law professors in Martínez and Winter's (2021a) sample at least somewhat agreed that there were predictable, feasible mechanisms through which the law could influence the long-term future, both in general and with regard to specific risks and via different areas of law. These main findings held true independent of demographic factors such as age, gender, political affiliation, and legal training, strongly suggesting that academic legal experts across the English-speaking world endorse to a surprisingly significant degree the normative and descriptive assumptions of legal longtermism. Investigating this question with regard to experts in other fields as well as a more general audience (e.g., cross-culturally) could give experimental legal philosophers (and legal actors

more generally) a better sense of whether people endorse legal longtermism and to what degree.

With regard to (b) why people hold those beliefs, philosophical-level XLJ is concerned with understanding both the *proximate cause* of longtermism-related beliefs (cognitively, what leads people to hold certain longtermist-related beliefs and/or engage in longtermist-related behavior) and *ultimate cause* of those beliefs (what evolutionary or adaptive forces gave rise to them in the first place).

For example, with respect to the normative assumption, there are many reasons to expect cognitively why people may assign a significantly higher legal-social discount rate to people in the far future than those who have thought more about this issue from a philosophical perspective (i.e., moral philosophers).¹³ Decades of work in the fields of behavioral economics and cognitive psychology have demonstrated that human judgment is prone to various sorts of cognitive biases, many of which are believed to influence our thinking about the long-term future (Yudkowsky, 2008; Ord, 2020; see also Schubert, Caviola, & Faber, 2019), including the following:

1. Present bias and hyperbolic discounting, the trend of overvaluing immediate rewards and undervaluing long-term consequences (e.g., Mischel & Ebbesen, 1970; O'Donoghue & Rabin, 1999; O'Donoghue & Rabin, 2015)
2. Scope insensitivity, the inability to value a problem with a multiplicative relationship to its size (e.g., Desvovages, Naughton, & Parsons, 1992; Kahneman, Ritov, Schkade, Sherman, & Varian, 1999; Slovic, 2010)
3. Diminishing marginal utility of life (Greene & Baron, 2001)
4. The availability heuristic, the tendency to heavily weigh judgments based on information that is available and/or can be readily recalled (Tversky & Kahneman, 1973; Schwarz et al., 1991; Gilovich, Griffin, & Kahneman, 2002).

Since work on cognitive biases suggests that experts are often just as susceptible to these biases as non-experts,¹⁴ it stands to reason that there may not be much of a legal expertise effect when assessing group differences in longtermist-related views. Indeed, as mentioned earlier, Martínez and Winter (2021a) found similar levels of endorsement for the normative premise of legal longtermism in legal experts as those observed in lay adults. On the other hand, it is plausible that proximate causes other than expertise could lead to revealed group differences, such as based on political affiliation. For example, recent political psychology literature has revealed ideological differences in the expanse of empathy (Waytz, Iyer, Young, & Graham, 2016), compassion, and moral circle (Waytz, Iyer, Young, Haidt, & Graham, 2019); political conservatives appear to expend their empathy toward more local targets (i.e., smaller, closer, more well-defined, and less encompassing social circles), whereas liberals tend to empathize with more global targets (i.e., larger, farther, less structured, and more encompassing social circles, including nonhumans). Insofar as future generations are considered more “global targets,” one would likewise expect ideological differences in the expanse of empathy and compassion toward them, particularly with regard to those in the far future as opposed to the near future.¹⁵

With regard to the ultimate cause of beliefs toward legal longtermism, there are also plausible evolutionary explanations for why humans would not have developed certain longtermist beliefs. For example, while it is relatively easy to imagine why it would have been useful for humans to develop an inclination toward protecting near-term future generations (e.g., to ensure that one's children and grandchildren survive and flourish), it seems more difficult to imagine how an impulse toward protecting long-term future generations would have developed. After all, there would never be any opportunity in one's lifetime to directly act on this impulse nor any available feedback mechanism to observe the results of acting on this impulse, as far-future descendants would come into existence long after the end of that lifetime. Furthermore, while reciprocal altruism is observed among individuals and groups that are unrelated and sometimes physically distant from one another (e.g., Trivers, 1971), the mechanisms underlying such behavior seem unlikely to lead to prosocial tendencies toward future generations, given the impracticability of reciprocally cooperating with generations not living at the same time.

That said, with respect to both the proximate and ultimate cause of legal longtermist beliefs, many of these hypotheses and predictions have yet to be tested experimentally, leaving them ripe for exploration by XLJ researchers.¹⁶

10.2.2 Doctrinal-Level XLJ

Step 1 of doctrinal-level XLJ involves investigating the extent to which longtermism maps onto people's understanding of law, the legal system, and legally relevant concepts. As discussed in Tobia (2020a), legally relevant concepts can take various forms and may or may not have both a legal and an ordinary language counterpart. These forms include the following:

1. Legal concepts such as "promissory estoppel" and "subject matter jurisdiction," which do not have an ordinary language counterpart
2. Hybrid concepts such as "rights," "duties," and "person," which exist as legal concepts and are also used by laypeople in an everyday, non-legalistic context
3. Ordinary concepts such as "January," "dollar," and "vegetable," which do not have an explicitly legal counterpart but may nonetheless be legally relevant (see, e.g., *Nix v. Hedden*, 1893).

XLJ is concerned with each of these types of concepts, insofar as one can advance legal arguments on the basis of their experimental investigation.¹⁷

With regard to legal concepts without an explicit ordinary language counterpart, for example, a recent doctrinal-level project investigated to what degree law professors' understanding of the concept of standing (*locus standi*) extended to future generations.¹⁸ In a set of common-law-trained law professors from around the world, Martínez and Winter (2021) found that a slight majority of respondents leaned toward or accepted the proposition that there was a reasonable legal basis for granting standing to humans living in the near future (understood as up to 100 years from now), while slightly more than one-third of participants leaned toward or accepted the proposition for humans

living in the far future (understood as over 100 years from now). The figure for humans living in the near future was slightly higher than that for non-human animals and sentient artificial intelligence but lower than that for groups such as corporations, unions, and, perhaps surprisingly, the environment (understood as rivers, trees, or nature itself).

For *locus standi* and other specialized legal concepts that do not have an explicit ordinary language counterpart and are unfamiliar to those without legal training, surveying legal experts to the exclusion of laypeople (i.e., those unfamiliar with legal doctrine) may make more sense.¹⁹ However, for classes of legally relevant concepts that do have an ordinary language counterpart familiar to laypeople—such as hybrid concepts like personhood, rights, and duties—XLJ is concerned with understanding how ordinary people understand those concepts in relation to longtermism, as well as to what degree this understanding differs from that of legal experts. For example, in Martínez and Winter's (2021b) survey of over 1000 US adults, 64.09 percent of participants considered at least some subset of humans living in the near future to be persons, and 61.75 percent considered at least some subset of humans living in the far future to be persons. Similar results were observed in Martínez and Winter's (in press a) cross-cultural survey of lay adults in ten different countries. In both studies, the percentage was higher than that observed in Martínez and Tobia's (2021) survey of over 500 law professors, in which just over 50 percent of participants considered at least some subset of humans living in the near future to be persons, and fewer than 50 percent considered at least some subset of humans living in the far future to be persons.

As will be further discussed in Section 10.3.2, from a legal perspective much of the emphasis on surveying lay intuitions—either in addition to or in lieu of those of legal experts—is based on the doctrine of ordinary meaning analysis, particularly prominent in the United States and other common-law jurisdictions but relevant to other jurisdictions as well, which states that words in a legal document should be interpreted according to their ordinary meaning.²⁰ In cases where ordinary meaning analysis does not apply, as well as in jurisdictions that do not have this doctrine or any equivalent to begin with, one might question the utility of surveying laypeople about legal concepts. However, as Tobia (2022) points out, most experimental jurisprudence studies, regardless of the jurisdiction, tend to emphasize laypeople and rarely use legal experts as subjects except in tandem with a lay sample, suggesting a potentially deeper reason for the utility of surveying laypeople.²¹

While one might expect a priori that similar cognitive (and evolutionary) factors would be at play in shaping how and to what extent longtermism fits into both expert and laypeople's understanding of legally relevant concepts, recent evidence suggests that those with legal training (including judges and even law students) interpret ordinary legal concepts (such as intentionality) differently from those without legal training (e.g., elite non-law students and the general public: Tobia, 2020b).

That said, it is unclear how legal training might affect people's judgments. In some cases one might expect legal training to result in a less longtermist-friendly interpretation of certain concepts.²² For example, when interpreting the concepts of "rights" and "duties," lawyers might be influenced by the fact that most jurisdictions

do not currently grant many legal rights to future generations, nor do they impose many legal obligations on current generations to protect future generations, whereas laypeople may simply interpret those concepts in a normative, jurisdiction-independent sense. With regard to personhood, although those with legal training in many respects have a more expansive interpretation of the concept of personhood that extends to corporations and other entities beyond “natural persons,” in other respects they have a narrower interpretation of the concept that is detached from the concept of “human,” hence explaining the discrepancy between the endorsement rates observed among lay participants (Martínez & Winter, 2021b) (with respect to considering future generations’ personhood) relative to expert participants (Martínez & Tobia, 2021).²³

10.2.3 Applied-Level XLJ

Step 1 of the third level of abstraction, which we refer to as applied-level XLJ, involves examining people’s intuitions about how more specific areas of law or even concrete legal mechanisms—such as constitutional provisions, congressional statutes, and agency regulations—could address more concrete long-term challenges, such as existential risks posed by artificial intelligence, extreme climate change, and synthetic biology. Note that while the use of the term *applied* might suggest that this level is dependent on the previous two levels, applied-level XLJ research can be undertaken either independently of or in tandem with research at the previous two levels.²⁴ Here, we discuss a few examples of applied-level XLJ, as well as their relationship to philosophical-level and doctrinal-level XLJ.

Whereas previously discussed XLJ research at the philosophical level found that the vast majority of law professors surveyed endorse the feasibility assumption of legal longtermism (i.e., that law can predictably and feasibly influence the long-term future), the applied-level research investigated their beliefs regarding how and in what ways the law can influence the long-term future (Martínez & Winter, 2021a). In a set of law professors ($n \approx 170$), subjects were asked whether they believed there were predictable, feasible mechanisms through which different areas of law could influence the long-term future (understood as at least 100 years from now) and very long-term future (understood as at least 1000 years from now). These questions represent applied-level research, as the different areas of law would be the means of influencing the long-term future. For each area of law surveyed, the majority of participants endorsed²⁵ the proposition that there were predictable, feasible mechanisms through which to influence the long-term future, though the percentage of those who agreed was significantly higher for some areas (environmental law, constitutional law, and property law) than for others (contract law and criminal law). Participants were also asked whether they believed there were predictable, feasible mechanisms through which law as a whole could influence the (very) long-term future with regard to different real-world issues. As with the areas of law, for each of the issues surveyed, the majority of participants at least somewhat agreed with the proposition that there were feasible, predictable legal mechanisms that could influence the long-term future, while the mean level of agreement was significantly higher for some areas

(e.g., climate change) than for others (e.g., artificial intelligence) (Martínez & Winter, 2021a).

In addition to investigating the long-term predictability and feasibility of general areas of law, two other approaches to applied-level research include investigating existing and/or proposed legal instruments with potential longtermist implications and assessing subjects' intuitions regarding (a) normatively, whether they would be in favor of implementing said instruments and/or (b) descriptively, how they would interpret the said instrument—once implemented—in various longtermism-relevant scenarios.

With regard to the first approach, on normative intuitions, a researcher might investigate—either directly, through survey methods, or indirectly, through a series of contrastive vignettes—expert and/or lay²⁶ intuitions regarding, for example, statutorily mandated budgets to reduce existential risk. Are most people in favor of implementing the said budgets? If so, how much of their jurisdiction's total spending do they think should be allocated to the said budgets (both overall and relative to other types of spending, such as military expenditures)? Are people's judgments influenced by cognitive factors such as scope neglect, thus causing those who might otherwise endorse longtermist legal protection to disregard the potential necessity or benefit of existential-risk-related legal mechanisms? These types of questions, as well as the methodological techniques to answer them, closely resemble those asked at the philosophical level, the main difference being that both the longtermist scenarios and the legal mechanisms related to those scenarios are much more specific at the applied level than at the philosophical level.

With regard to the second approach, a researcher might investigate people's intuitions regarding a constitutional provision stating that the government “should protect current and future generations against existential threats.” Does this constitutional provision require the government to implement budgets to reduce existential risks? Who has standing to sue on future generations' behalf if such protection is not provided? What counts as an existential threat according to this provision? How might one's interpretation of this provision differ for a constitutional provision compared to a statutory or regulatory provision? With regard to the latter question, pilot evidence from Martínez and Winter (2021) suggests that various sorts of hypothetical constitutional provisions (such as a commitment to spend 1 percent of GDP toward protection against existential risk, or a provision granting standing to future generations), are interpreted by legal experts as granting similar levels of protection to future generations. Further work is needed to confirm this finding as well as to investigate the other questions posed here.

It also remains unclear under what circumstances longtermists might prefer a precise or vaguely worded provision. On the one hand, for similar reasons discussed with regard to the philosophical and doctrinal level, one might expect a priori that more specific provisions (whether constitutional, statutory, or regulatory in nature) would tend to be interpreted more favorably toward future generations than would vaguer ones. First, if the judge (or other interpreter) does not endorse legal longtermism, then to the extent that they do not want to protect future generations, the language in a legal instrument must be sufficiently specific so as to ensure a

longtermism-friendly interpretation. Second, even if a judge does endorse legal longtermism, they may not decide in a way that is longtermism-friendly in cases where human reasoning is especially prone to faulty statistical intuitions or biases (cf. Schubert, Caviola, & Faber, 2019). This may be an issue specifically with regard to existential risks, which require reasoning based on very small probability scenarios of enormous magnitude and therefore might be influenced by factors such as scope neglect.

On the other hand, there are other reasons to expect that vague *standard*-like provisions might lead to broader protection of future generations in certain circumstances over more precisely drafted *rule*-like provisions. First, in cases where a judge wants to help future generations, a vague or abstract provision will allow the said judge to choose the broadest interpretation possible, whereas a precisely worded provision would constrain them to wording chosen by a particular legislator. Second, vague norms also allow for potentially increasingly longtermism-friendly interpretations over time, as future judges (and laypeople) might well extend their moral circle over time²⁷ and, by extension, be more likely to adopt a longtermism-friendly interpretation.²⁸

In terms of specific types of legal instruments and mechanisms, since constitutional provisions are generally more vague and broad than statutes, then insofar as vague provisions tend to lead to more longtermism-friendly interpretations, so too would one expect constitutional mechanisms to be more longtermism-friendly than statutory mechanisms, particularly since the former, *ceteris paribus*, constitutes higher and more powerful law.²⁹ Conversely, insofar as more specific provisions seem friendlier to future generations, then one would expect statutory mechanisms to likewise be friendlier toward future generations, as well. Moreover, depending on the jurisdiction, one would also expect in this regard that regulations, whose provisions are generally even more fleshed out than statutes and are often just as binding on the judge as legislation-passed statutes, to be longtermism-friendly, as well.³⁰

With regard to the second approach on descriptive intuitions, it seems reasonable to expect that there would be significant differences between experts and non-experts in their interpretations of specific longtermist legal mechanisms. Given the aims of applied-level XLJ, it seems reasonable here to focus on expert intuitions as opposed to lay intuitions (since the former will be more representative of the ultimate decision maker when the instrument is to be applied), unless there is reason to believe that expert and non-expert intuitions do not significantly deviate in certain cases at hand. Furthermore, given the large inter-jurisdictional variation in both the design and interpretation/application of legal instruments, it is important to ensure that findings of one jurisdiction are shown to be replicable in other jurisdictions before generalizing results.³¹

Taking into account these and other relevant methodological considerations, applied-level XLJ would help provide insight into which sorts of prospective legal instruments are most likely to promote longtermist goals and which are more likely to be ineffectual or harmful toward longtermist goals. We discuss these normative implications (as well as those with regard to the philosophical and doctrinal levels) in the next section.

10.3 Normative Implications of Experimental Longtermist Jurisprudence

While Section 10.2 focused on the empirical aims of experimental longtermist jurisprudence research at three levels of abstraction, here we discuss whether and how those empirical findings might inform normative discussions. This section is likewise split up into three subsections and covers philosophical implications (10.3.1), legal implications (10.3.2), and policy implications (10.3.3). Note that while each of these types of implications generally maps on to a distinct level of abstraction covered in the previous section, there is also a significant degree of overlap, which we acknowledge when relevant.

10.3.1 Philosophical Implications

Philosophically, XLJ is concerned with determining whether and to what extent future generations ought to be provided legal protection (independent of existing legal doctrine). Here, we outline three general approaches similar to other forms of experimental philosophy: (a) the *if-then* approach, (b) the *debunking* approach(es), and (c) the *pluralism* approach.³² Note that these relate mostly to philosophical-level XLJ, though the doctrinal level and applied level may often be relevant, as well.

The if-then approach is the most straightforward and essentially states that if relevant participants consistently make a judgment in favor of a particular moral (or legal) claim, then that claim has prima facie normative weight. Note that “relevant participants” here is potentially open to interpretation and may refer to laypeople or certain types of experts, depending on one’s normative lens. Here we may distinguish between *democratic if-then* approaches to XLJ, which involve drawing normative inferences based on the judgments of laypeople, and *technocratic if-then* approaches, which involve drawing normative inferences based on the judgments of experts. With regard to both democratic and technocratic if-then approaches, the normative import provided by participants’ judgments alone seems quite limited due to classic is-ought concerns (i.e., just because most people believe X does not mean that most people should believe X, nor does it mean that X is true³³), and the said judgments will never be enough on their own to deliver an “all-things-considered” normative conclusion, even in cases where all actors agree.³⁴ However, if one either assumes or makes additional arguments in favor of the reliability and trustworthiness of laypeople’s or experts’ intuitions in relevant contexts, then the if-then approach may help deliver valuable philosophical insights.

While the limitations of the if-then approach seem to apply to its use in justifying the normative premise of XLJ (all consequences matter equally), this approach (particularly the technocratic if-then approach) seems more promising as a way of justifying the empirical premises of legal longtermism, particularly the feasibility assumption. For example, to the extent that legal academics are experts on the potential

long-term effects of law,³⁵ it follows that their endorsement of the claim that there are feasible, predictable mechanisms through which the law can influence the long-term future would strengthen the same empirical premise underlying legal longtermism (i.e. that there are feasible, predictable mechanisms through which the law can influence the long-term future), which in turn would provide some evidentiary and normative weight to legal longtermism.

The second set of approaches to drawing normative philosophical inferences from XLJ is known as *debunking*. Whereas the if-then approach assigns normative weight to relevant participant judgments outright, debunking approaches argue against assigning normative weight to judgments that are unreliable. Here again, “unreliable” can be open to interpretation and is essentially a catchall term for morally irrelevant factors or processes. According to one prevalent version of this approach, often referred to as *cognitive debunking*,³⁶ a judgment is considered to be unreliable if the underlying psychological process giving rise to that judgment—or, to use the terminology used in Section 10.2, the proximate cause of that judgment—does not reliably “get to the truth” (Wedgwood, 2007) or “track the truth” (Andow, 2016), or it cannot be classified as a “truth-tracking process” (Greene, 2013, 2014; Winter, in press).

The debunking approach is perhaps the most formalized of the philosophical approaches covered in this section and can be represented as follows:

- (P1) Judgment p is the output of a psychological process that is substantially influenced by factor F. (Empirical premise)
- (P2) If a judgment is the output of a psychological process that is substantially influenced by factor F, then it is pro tanto unreliable. (Normative premise)
- (C) Judgment p is pro tanto unreliable.

With regard to legal longtermism, for example, let us say that p is the judgment that future generations must not be protected under the law to the same degree as the current generation, and F is one of the cognitive biases laid out in Section 10.2. According to the debunking approach, to the extent that the underlying cognitive processes of this judgment are influenced by these cognitive biases, this judgment is unreliable.

While the debunking approach is often used to lower the normative weight of one judgment, one might instead decide to assign different normative weight to the judgments of laypeople and experts depending on the context. This approach, sometimes referred to as pluralism (Earp et al., 2021), might particularly apply at the doctrinal level, where laypeople and legal experts might both have reliable but competing judgments of longtermist-relevant ordinary legal concepts, such as *person(hood)*. In these cases, one might assign normative weight to the lay judgments in cases where the word is used in its “ordinary” sense and to expert judgments in cases where the word appears to be a “term of art.” This distinction will be further explored in the next section as we discuss the legal implications of XLJ.

10.3.2 Legal Implications

In addition to providing philosophical arguments for why and to what extent future generations ought to be provided legal protection in an ideal legal system, XLJ might also be used to advance *legal* arguments regarding the extent to which future generations ought to be provided legal protection according to the doctrines of current legal systems. Here we discuss how doctrinal-level XLJ might be utilized to advance such normative legal arguments, using as illustration two principles of legal interpretation: ordinary meaning and terms of art.

Legal interpretation is ubiquitous to legal argumentation and decision-making, and in many jurisdictions, ordinary meaning analysis has been referred to as “the most fundamental principle” of legal interpretation (Slocum, 2015).³⁷ According to the ordinary meaning rule, words in a statute (see, e.g., *Moskal v. United States*, 1990; *United States v. Turkette*, 1981; *Richards v. United States*, 1962), treaty (see, e.g., Vienna Convention on the Law of Treaties art. 31, 1969; Slocum & Wong, 2021), contract (*California Civil Code*, 2018; *Jowett, Inc. v. United States*, 2000; *Harris v. Dep’t of Veterans Affairs*, 2015) or other legal document should generally be interpreted according to their ordinary meaning or usage (as opposed to their technical definition, for example).³⁸ Although there is debate as to what “ordinary meaning” itself means, most jurists seem to agree that it, to some extent, encompasses how a typical or reasonable person generally understands and uses a given word or concept (see, e.g., Tobia, 2020b; Lee & Mouritsen, 2018; Klapper, Schmidt, & Tarantola, 2020).³⁹ Since one of the goals of XLJ at the doctrinal level is to discover how the typical person understands and interprets longtermism-relevant words and concepts, the findings of doctrinal-level XLJ could plausibly be used to advance ordinary-meaning-related legal arguments (see, e.g., Martínez & Winter, 2022, investigating the ordinary meaning of existential risk).

Of course, not all cases of interpretation involve ordinary meaning analysis, particularly in cases involving terms of art—that is, words that have a particular meaning in a field (such as law, science, or business) that deviates from that word’s ordinary meaning or usage (see, e.g., *Corning Glass Works v. Brennan*, 1974; *Housey Pharm., Inc. v. Astrazeneca U.K. Ltd.*, 2004; *Occidental Life Ins. Co. of Cal. v. United States*, 1965). In such cases, where the terms are to be given their *technical* as opposed to *ordinary* meaning (e.g., Frankfurter, 1947), XLJ could be useful in terms of (a) identifying terms of art (cf. *Nix v. Hedden*, 1893) and (b) interpreting terms of art, once identified. For example, in cases where it is unclear whether a given longtermism-relevant word is a term of art, one could compare the interpretations of ordinary people and legal experts (or experts in the relevant field) to verify whether the two interpretations tend to deviate significantly from one another. If they do, one could advance a legal argument in favor of interpreting the word as a term of art as opposed to performing ordinary meaning analysis.⁴⁰ Furthermore, in cases where a longtermism-relevant term has already been identified as a term of art, one could use the judgments of experts gathered through XLJ research to advance legal arguments in favor of using those judgments to interpret the term of art in question.

Note that these approaches to advancing legal arguments are similar but not identical to the pluralistic approaches to advancing philosophical arguments identified in the previous section. While the pluralistic approach could be used to identify

how the law should conceptualize longtermism-relevant concepts, independent of current legal doctrine, the approaches outlined here could be used to determine and convince legal actors—such as judges—how longtermism-relevant concepts should be interpreted according to existing legal doctrine. In other words, whereas the normative legal arguments discussed in this section could plausibly be used in the courtroom to address questions of law,⁴¹ many of the normative philosophical arguments could not.

10.3.3 Policy Implications

In addition to the aforementioned philosophical and legal normative implications, at the policy level XLJ is concerned with determining which legal mechanisms can and ought to be implemented so as to provide the appropriate level of legal protection to future generations. Two ways of doing so include (a) identifying the most implementable longtermist policies and legal instruments in a given jurisdiction and (b) identifying legal instruments that, once implemented, would most effectively protect the interests of future generations. Here we discuss each of these approaches in turn.

The first approach to drawing policy implications involves using XLJ findings—particularly at the applied level, though to some extent at the philosophical and doctrinal levels, as well—to identify longtermism-relevant policies that would be more feasibly implemented. For example, suppose that an applied-level XLJ research project finds that the vast majority of citizens in a democratic jurisdiction are in favor of implementing mandatory federal budgets for protection against existential threats but are not in favor of granting future generations personhood status. To the extent that legislative decision makers are receptive to the popular will of their constituents, this would imply that implementing mandatory federal budgets for existential threats would be more feasible than granting future generations personhood status, and further that, from a legal longtermist perspective, *ceteris paribus*, one should prefer to attempt to implement the former as opposed to the latter.⁴²

The second approach to drawing policy implications involves using XLJ findings to identify legal instruments that, once implemented, would be most favorably interpreted from a longtermist perspective. For example, suppose that an applied-level XLJ research project investigating two versions of a longtermism-relevant provision finds that the vast majority of legal experts interpret version A of the provision much more favorably toward future generations in most relevant scenarios than version B of the provision. Insofar as legal experts' interpretations are reflective of that of the judge tasked with applying the provision in a longtermism-relevant case, it would follow that from a longtermist perspective, *ceteris paribus*, one should prefer to implement version A of the provision as opposed to version B.

Although these might seem like overly specific policy implications, at a general level one might also be able to draw policy implications regarding the relative feasibility and desirability of (a) broad versus narrow provisions; (b) constitutions versus statutes, regulations, and other legal instruments; and (c) rights versus privileges, duties, and other legal concepts within longtermism-relevant legal instruments. For example, as briefly discussed in Section 10.2.3, there are *a priori* reasons to believe that both broad

standards and narrow rules might be preferable from a longtermist perspective in various scenarios. Applied-level XLJ research could test these predictions and provide general policy guidance to those hoping to implement longtermism-relevant legal mechanisms.

Note that while this approach of drawing normative policy implications from interpretive intuitions may seem superficially similar to the approach for drawing normative *legal* implications from such intuitions, the two differ in subtle yet important respects. For example, consider a proposed longtermism-relevant piece of legislation that is found to be interpreted by the majority of laypeople to offer legal protection to future generations, while the majority of legal experts interpret it as not offering legal protection to future generations. In this case, one might argue from a longtermist policy perspective that it should not be implemented due to the fact that since legal experts interpret it as not offering legal protection to future generations, that the deciding judge will likewise interpret it as not offering legal protection to future generations. However, supposing that this legislation were in fact implemented, then one might make the normative *legal* argument that the ordinary meaning of the text implies legal protection to future generations, since the majority of laypeople—whose intuitions are arguably a better reflection of ordinary meaning than those of a legal expert—interpret the statute in that manner and that consequently the judge similarly ought to interpret the text in this manner.⁴³

10.4 Conclusion

This chapter has presented a new research area within the burgeoning field of experimental jurisprudence aimed at informing philosophical, legal, and political debates related to the long-term future. As alluded to in Section 10.1, longtermism is a comparatively new research area in the realm of philosophy and has even more recently made its way into the realm of law and jurisprudence. Whereas many areas of study within experimental philosophy and experimental jurisprudence focus on issues of long-standing discussion within general philosophical literature, XLJ offers the rare opportunity within these fields to apply cutting-edge methodological approaches to cutting-edge substantive issues relevant to philosophers, lawyers, and policy makers. Indeed, XLJ research has already shown that the vast majority of legal scholars consider the protection of the long-term future of utmost importance (Martínez & Winter, 2021a).

At the same time, much of the early work in XLJ, though promising, has only scratched the surface in terms of answering both the empirical questions posed in Section 10.2 and the normative debates of Section 10.3, and we encourage both experimentalists familiar with the methods outlined in Section 10.1.3 as well as theorists familiar with the substantive questions outlined in Sections 10.1.1 and 10.1.2 to work on advancing XLJ research at each of the three levels of abstraction introduced in this chapter. Furthermore, this levels of abstraction framework not only serves to lay out the diverse set of empirical and normative aims within XLJ but also contributes to a more general conceptualization within experimental and theoretical jurisprudence,

and we invite researchers within the broader sphere of legal philosophy to incorporate this framework into future projects when useful.

Acknowledgments

For comments, discussion, and critique, we are grateful to Kevin Tobia, Sean Richardson, and Piotr Bystranowski, as well as our editor and reviewer. We also want to thank Suzanne Van Arsdale in particular for help with the manuscript and many useful suggestions.

Notes

- 1 Note that experimental jurisprudence may also be referred to as *experimental legal philosophy*, though for the sake of simplicity we will use the terminology experimental jurisprudence throughout the chapter.
- 2 However, there is not yet a widely accepted definition of longtermism (see, e.g., MacAskill, 2019). Also note that there are several different versions of longtermism. The version that we lay out here is most similar to “weak longtermism,” which holds merely that we should be particularly concerned with ensuring that the long-run future goes well, as opposed to “strong longtermism,” which holds that impacts on the long-run future are the most important feature of our actions. For more information on further distinctions within strong longtermism, see Greaves and MacAskill (2021). For more information on the philosophical foundations of longtermism more generally, see Parfit (1984); Beckstead (2013, 2019); and Greaves and MacAskill (2021).
- 3 Note that our presentation of the case for longtermism deviates slightly from the presentation in the formal longtermist literature in that (a) we provide three assumptions instead of two (the feasibility argument is not presented as an “assumption” in longtermist literature), and (b) we do not explicitly discuss any of the objections to longtermism. These deviations are largely for the purpose of brevity and readability; for a more thorough discussion of the objections, see Greaves and MacAskill (2021) and Tarsney (2020). For a concise overview of these objections that is tailored to a legal audience, see Winter et al. (2021).
- 4 An alternative framing of this assumption would be to state that (a) all consequences matter equally insofar as they affect welfare, and (b) the welfare of those living in the future matters equally as the welfare of those living today. Note, however, that the use of the term “consequences” does not necessarily imply that longtermism is an inherently consequentialist theory, as consequentialism is just one approach to justifying longtermism. For example, one might alternatively argue (from a deontological perspective) that we owe a duty to future generations, independent of what a consequentialist/utilitarian calculus might demand, or (from a virtue ethics perspective) that it is a virtue to act in a way that protects future generations by exercising patience, self-discipline, benevolence, and taking responsibility for our actions (Gaba, 1999, pp. 283–7; cf. also Ord, 2020). One might also conceivably value future generations from a purely aesthetic or intellectual achievement standpoint (Todd, 2017). Our apparent focus on consequences in

discussing the assumptions of longtermism stems from the fact that most non-consequentialist ethical theories maintain that, *ceteris paribus*, consequences matter to some degree.

- 5 While it may still seem counterintuitive that we should not adopt a “social discount rate” (as we do for money, for example), this view seems to be shared by most, if not all, moral philosophers that have written about the issue (Parfit, 1984; Cowen & Parfit, 1992; Broome, 1994; Mogensen, 2019), as well as by many theoretical economists (38 percent according to a survey performed by Drupp, Freeman, Groom, & Nesje, 2018). As we outline later on, XLJ research has uncovered that legal scholars also seem to broadly share this view (or at least a much smaller discount rate than is currently being applied).
- 6 Note that although the paragraph argues mostly from the perspective of humans, “individuals” here can refer not only to humans but to potentially all other forms of sentience, as well (e.g., non-human animals and sentient artificial intelligence [assuming its existence]). Indeed, there appear to be very compelling arguments for including this broader set of sentient beings within a moral calculus, and, by extension, within the definition of individuals as presented above (Bentham, 1789; Singer, 1973; Gruen, 2017). That said, independent of how one defines individuals here, the associated arguments—both for this assumption and for longtermism as a whole—would proceed similarly as presented in the main body text.
- 7 The usual threats of extinction faced by species include environmental, demographic, and genetic factors (Benson et al., 2016).
- 8 See generally the emerging field of persistence studies, for example, Giuliano and Nunn (2021) with further references therein; see also Pirie (2021).
- 9 Note that in terms of justifying longtermism, the extent to which premise three (feasible and predictable) must be true arguably depends in part on the extent to which premise two (future is vast in size) turns out to be true (and vice versa). For example, the greater in size the future turns out to be, the more one can be confident in longtermism despite less feasible, predictable ways of influencing the long-term future. Conversely, the more feasible and predictable one believes it is to influence the long-term future, the smaller the number of future individuals must be for one to conclude that longtermism is true. To some extent, this is also the case with regard to premise one (all consequences matter equally); even if, in spite of the reasons laid out above, one does not believe the future is worth valuing *to the same degree* as the present, and thus discounts the value of future generations as discussed in note 5, one could still conclude that longtermism is true depending on the expected size of future generations, feasibility and predictability of influence, and degree of discounting. Note that the interrelation of these premises/assumptions influences not only confidence in longtermism but also confidence in weak longtermism versus strong longtermism discussed in note 2; that is, depending on such calculations, one might not only conclude that we should be “particularly” but rather “primarily” concerned with ensuring that the long-run future goes well (cf. Greaves & MacAskill, 2021).
- 10 As with philosophical longtermism, there are several possible versions of legal longtermism. For example, analogous to *strong philosophical longtermism*, we may define *strong legal longtermism* as the view that “the primary determinant of the value of a legal mechanism is the effect of that mechanism on the far future.”
- 11 For an overview of experimental jurisprudence, see Tobia (2022). For an overview of experimental philosophy more generally, see Knobe et al. (2012).

- 12 As we clarify further, note that “people” here can refer not only to laypeople but also to legal experts, within or across different jurisdictions depending on the study.
- 13 See *supra* note 5. See also Greaves (2017) for a survey of discounting in public policy, including a survey of the arguments for and against a positive rate of pure time preference. Among moral philosophers, a zero rate of pure time preference is endorsed by, among others, Broome (2008), Buchholz and Schumacher (2010), Cline (1992), Cowen and Parfit (1992), Dasgupta (2008), Dietz, Hepburn, and Hope (2008), Gollier (2012), Harrod (1948), Pigou (1932), Ramsey (1928), Sidgwick (1907), Solow (1974), and Stern (2008). Other related philosophical works include Cowen and Parfit (1992), Mogensen (2019), and Parfit (1984). In a survey of experts on social discounting, 38 percent accepted a zero rate of pure time preference (Drupp et al., 2018).
- 14 For a comprehensive list of relevant studies on expert intuition, see Guthrie, Rachlinski, and Wistrich (2000); see also Winter (2020, p. 249).
- 15 That said, the results from Martínez and Winter (2021) indicate that conservatives and liberals alike endorse greater levels of legal protection for future generations than those granted to future generations under current legal institutions, suggesting that future generations may be within the moral circle of both liberals and conservatives.
- 16 Note that with regard to the ultimate cause, computational approaches—such as game-theoretic modeling—may be a more promising means for uncovering legal longtermism-relevant insight than experimental approaches. In other words, this might be better suited for the realm of what we might label as *computational longtermist jurisprudence* as opposed to experimental longtermist jurisprudence, though the two approaches are, of course, by no means mutually exclusive.
- 17 Connecting the second level more explicitly with the first level, one might also ask to what extent people’s endorsement (or lack thereof) of legal longtermism is a result of or relates to their conception of ordinary legal concepts.
- 18 Standing (*locus standi*) in many jurisdictions refers to the legal right of one to bring forth a lawsuit.
- 19 As will be discussed further in Section 10.3, note that surveying laypeople on issues *underlying* specialized concepts (e.g., should future generations be represented in court) may still make sense at the philosophical or applied level as opposed to the doctrinal level.
- 20 Examples of jurisdictions that explicitly employ some version of ordinary meaning analysis include Australia (e.g., *Electricity Generation Corporation v Woodside Energy*, 2014), the United Kingdom (*River Wear Commissioners v. Adamson*, 1877), South Africa (*Venter v. R*, 1907; Carney, 2016), the United States (see generally Slocum, 2015), and Singapore (Interpretation Act Sec. 9A, 1993), as well as international law (Vienna Convention on the Law of Treaties art. 31, 1969). Ordinary meaning has also been found to be relevant in civil-code jurisdictions, as well, including Argentina, Finland, France, Germany, Italy, Poland, and Sweden (see generally McCormick & Summers, 2016).
- 21 Alternatively, one may argue that the use of lay people reflects a bias for convenience, as it is much easier to recruit lay subjects than legal professionals. We further discuss these issues in Section 10.3, *infra*.
- 22 Compare this to the effect of philosophical training on people’s judgments of philosophical-based legal longtermism, discussed in note 5 and accompanying text.

- 23 For example, whereas laypeople may be more inclined to associate the concept of “person” with the concept of “human,” which would presumably lead to a relatively high degree of endorsement of the view of future humans as “persons,” lawyers may be more reticent to do so given that (a) future humans are not commonly thought of as legal persons, and (b) it might seem counterintuitive that future humans could do things that legal persons are commonly thought to be able to do (e.g., enter into contracts, sue and be sued, own property, etc.). On the other hand, it is also counterintuitive that corporations and other juridical persons could do these things, so perhaps the latter may not be so much of an impediment to a lawyer’s ability to consider future generations as legal persons (just as it may not impede consideration of non-human animals, for example), hence explaining why a slight majority still endorsed personhood for humans living in the near future.
- 24 In practice, however, research that investigates more concrete legal mechanisms is most likely carried out for instrumental purposes—that is, not for the purposes of determining (or even arguing) whether longtermism fits or ought to fit the goals of the current legal system, but rather how a longtermism-friendly policy maker might design legal instruments that will ultimately be interpreted in a longtermism-friendly way.
- 25 For each area of law, participants were asked, on a scale of 1–7 (1 being “strongly disagree, 4 being “neutral, and 7 being “strongly agree”), to rate their level of agreement with the following statement: “There are feasible, predictable mechanisms through which the law can influence the long-term future (understood as at least 100 years from now) via . . .” followed by the relevant area of law. For each area, the majority of participants responded 5 or higher (indicating at least “somewhat agree”).
- 26 For further discussion on the relative normative implications of using expert versus lay intuitions in this context, see Section 10.3.
- 27 Anthis & Paez (2021); Singer ([1981] 2011).
- 28 For example, the eighth amendment to the US Constitution (1791) prohibits “cruel and unusual punishment.” In the mid-twentieth century judges began to interpret this phrase as an “evolving standard” (*Trop v. Dulles*, 1958) with more progressive judges eventually interpreting it as including capital punishment despite the original interpretation of this provision not including capital punishment (*Furman v. Georgia*, 1972; Stinneford, 2019).
- 29 On the other hand, constitutions are surprisingly short-lived; Elkins, Ginsburg, and Melton (2007) found that (a) the average lifespan of a constitution is just seventeen years, and (b) the probability of a constitution lasting at least fifty years is just 19 percent. Although systematic data with regard to statutes appears unavailable, there are many cases of legislation outlasting several iterations of a nation’s constitution (e.g., the German Penal Code of 1871: Mueller, 1961).
- 30 Other potential reasons dictating in favor of regulatory mechanisms (particularly in the United States) are that they are more easily passed than statutes, can be updated more easily based on new information, and tend to be implemented by those with expert knowledge of the material in question.
- 31 It also seems likely that many of the legal mechanisms evaluated at the applied level are jurisdiction-specific; accordingly, in some cases it may make less sense to attempt to replicate some of the exact findings of one experiment in another jurisdiction and instead investigate these differences to understand what makes a particular jurisdiction have more longtermism-friendly interpretations.
- 32 Note that the labels for these approaches are not yet very well-established. We have adapted the terminology from Earp, Lewis, Dranseika, and Hannikainen (2021)

- regarding approaches to drawing normative implications in experimental bioethics, as we find this to be a useful taxonomy.
- 33 Despite the is-ought concerns, many have argued that moral philosophers are, in fact, moral experts (Singer, 1972).
 - 34 Although cases where all potentially relevant participants converge toward one view might provide normative weight toward that view, one would still need to rule out the possibility that the said actors are all systematically biased (see debunking approach later).
 - 35 It remains an open question to what extent legal academics (or legal experts in general) are the relevant experts here, as opposed to, say, superforecasters (cf. Tetlock & Gardner, 2016).
 - 36 Note that cognitive debunking is just one type of debunking approach. Another common debunking approach is evolutionary debunking, which involves assessing whether the evolutionary process that gave rise to a particular belief or judgment is unreliable (see generally Mogensen, 2016).
 - 37 See *supra* note 19.
 - 38 For an overview of the ordinary meaning analysis doctrine, see generally Slocum (2015).
 - 39 For example, in *Addison v. Holly Hill Fruit Products, Inc.* (1944), the court stated, “Legislation when not expressed in technical terms is addressed to the common run of men and is therefore to be understood according to the sense of the thing, as the ordinary man has a right to rely on ordinary words addressed to him” (p. 618).
 - 40 Note that at the philosophical level, one could make an argument for why the law *should* use the lay definition (independent of ordinary meaning analysis or any other doctrine) using the democratic if-then approach. Conversely, one could use a technocratic if-then approach to argue in favor of the interpretation of a legal expert (even for cases currently governed by ordinary meaning analysis).
 - 41 Note that legal interpretation is generally considered a question of law. For example, with regard to the United States, see, e.g., *United States v. Moore* (2009); *United States v. Shafer* (2009).
 - 42 Of course, the assumption that lawmakers are receptive to the interests of the public is not always warranted. For example, in the United States, the majority of Americans believe that the government has a responsibility to provide health care for all (Pew, 2020), yet the United States famously does not provide health care to all of its citizens.
 - 43 This hypothetical scenario, of course, assumes that ordinary meaning analysis would apply to this case, and that the relevant sections of the text would not be interpreted as terms of art (in which case the relevant experts’ intuitions would apply to any normative legal-interpretive argument as well).

References

- ACE Project (2021). *Comparative data*. <https://aceproject.org/epic-en>.
- Adams, F. C., & Laughlin, G. (1997). A dying universe: The long-term fate and evolution of astrophysical objects. *Reviews of Modern Physics*, 69(2), 337–72. <https://link.aps.org/doi/10.1103/RevModPhys.69.337>.
- Andow, J. (2016, June 7). Reliable but not home free? What framing effects mean for moral intuitions. *Philosophical Psychology*, 29(6), 904–11. <https://doi.org/10.1080/09515089.2016.1168794>.
- Anthis, J. R., & Paez, E. (2021, June). Moral circle expansion: A promising strategy to impact the far future. *Futures*, 130: 102756. <https://doi.org/10.1016/j.futures.2021.102756>.

- Araújo, R., & Koessler, L. (2021, September 30). *The rise of the constitutional protection of future generations* (Legal Priorities Project Working Paper Series No. 7-2021). Legal Priorities Project. <https://www.legalpriorities.org/research/constitutional-protection-future-generations.html>.
- Barnosky, A. D., Matzke, N., Tomiya, S., Wogan, G. O., Swartz, B., Quental, T. B., Ferrer, E. A. (2011, March 2). Has the earth's sixth mass extinction already arrived? *Nature*, 471(7336), 51–57. <https://doi.org/10.1038/nature09678>.
- Beckstead, N. (2013, May). *On the overwhelming importance of shaping the far future* [Unpublished doctoral dissertation]. Rutgers University. <https://doi.org/10.7282/T35M649T>.
- Beckstead, N. (2014, June 22). *Will we eventually be able to colonize other stars? Notes from a preliminary review*. Future of Humanity Institute, University of Oxford. <https://www.fhi.ox.ac.uk/will-we-eventually-be-able-to-colonize-other-stars-notes-from-a-preliminary-review/>.
- Beckstead, N. (2019). A brief argument for the overwhelming importance of shaping the far future. In H. Greaves & T. Pummer (Eds.), *Effective altruism: Philosophical issues*, (pp. 80–98). Oxford University Press. <https://doi.org/10.1093/oso/9780198841364.001.0001>.
- Beitz, C. R. (1983, October). Cosmopolitan ideals and national sentiment. *The Journal of Philosophy*, 80(10), 591–600. <https://doi.org/10.2307/2026155>.
- Benson, J. F., Mahoney, P. J., Sikich, J. A., Serieys, L. E., Pollinger, J. P., Ernest, H. B., & Riley, S. P. (2016, August 31). Interactions between demography, genetics, and landscape connectivity increase extinction probability for a small population of large carnivores in a major metropolitan area. *Proceedings of the Royal Society B: Biological Sciences*, 283(1837), 20160957. <https://doi.org/10.1098/rspb.2016.0957>.
- Bentham, J. (1789), *An introduction to the principles of morals and legislation*. Clarendon Press. <http://dx.doi.org/10.1093/oseo/instance.00077240>.
- Berman, H. J. (1985), *Law and revolution: The formation of the western legal tradition*. Cambridge, MA: Harvard University Press.
- Broome, J. (1994, April). Discounting the future. *Philosophy & Public Affairs*, 23(2), 128–56. <https://doi.org/10.1111/j.1088-4963.1994.tb00008.x>.
- Broome, J. (2008, June). The ethics of climate change: Pay now or pay more later?. *Scientific American*, 298(6), 96–102. <https://www.scientificamerican.com/article/the-ethics-of-climate-change/>.
- Buchholz, W., & Schumacher, J. (2010, September). Discounting and welfare analysis over time: Choosing the η . *European Journal of Political Economy*, 26(3), 372–85. <https://doi.org/10.1016/j.ejpoleco.2009.11.011>.
- Caldeira, K., & Kasting, J. F. (1992, December 31). The life span of the biosphere revisited. *Nature*, 360(6406), 721–3. <https://doi.org/10.1038/360721a0>.
- California Civil Code, Part 2. Contracts, § 1549–1701 (2018). <https://leginfo.ca.gov/faces/codesTOCSelected.xhtml?tocCode=CIV>.
- Carney, T. R. (2016, May 31). Using frames to determine ordinary meaning in court cases: the case of “plant” and “vermin.” *Stellenbosch Papers in Linguistics*, 45, 31–48. <https://doi.org/10.5774/45-0-209>.
- Ceballos, G., Ehrlich, P. R., Barnosky, A. D., García, A., Pringle, R. M., & Palmer, T. M. (2015, June 19). Accelerated modern human-induced species losses: Entering the sixth mass extinction. *Science Advances*, 1(5), e1400253. <https://doi.org/10.1126/sciadv.1400253>.
- Cheng, L., Rosett, A., & Woo, M. (Eds.). (2003). *East Asian law: Universal norms and local cultures*. London & New York: Routledge.

- Cline, W. R. (1992). *The economics of global warming*. Washington, DC: Institute for International Economics.
- Corning Glass Works v. Brennan, 417 U.S. 188 (1974).
- Cowen, T., & Parfit, D. (1992). Against the social discount rate. In P. Laslett & J. S. Fishkin (Eds.), *Justice between age groups and generations* (pp. 144–61). New Haven: Yale University Press.
- Dasgupta, P. (2008, September 4). Discounting climate change. *Journal of Risk & Uncertainty*, 37(2), 141–69. <https://doi.org/10.1007/s11166-008-9049-6>.
- Desvousges, W. H., Naughton, M. C., & Parsons, G. R. (1992, March). Benefit transfer: Conceptual problems in estimating water quality benefits using existing studies. *Water Resources Research*, 28(3), 675–683. <https://doi.org/10.1029/91WR02592>.
- Dietz, S., Hepburn, C., & Hope, C. (2008, March 28). Discounting and climate change: a non-marginal policy choice [Unpublished manuscript]. https://www.researchgate.net/publication/254809091_Discounting_and_climate_change_a_non-marginal_policy_choice.
- Drupp, M. A., Freeman, M. C., Groom, B., & Nesje, F. (2018, November). Discounting disentangled. *American Economic Journal: Economic Policy*, 10(4), 109–34.
- Earp, B. D., Lewis, J., Dranseika, V., & Hannikainen, I. (2021). Experimental philosophical bioethics and normative inference. *Theoretical Medicine & Bioethics*. <https://doi.org/10.1007/s11017-021-09546-z>.
- Electricity Generation Corporation v. Woodside Energy Ltd., 7 HCA (2014).
- Elkins, Z., Ginsburg, T., & Melton, J. (2007, December 26). *The lifespan of written constitutions* [Unpublished manuscript]. College of Law & Department of Political Science, University of Illinois. https://jenni.uchicago.edu/WJP/Vienna_2008/Ginsburg-Lifespans-California.pdf.
- Foote, M., & Raup, D. M. (1996). Fossil preservation and the stratigraphic ranges of taxa. *Paleobiology*, 22(2), 121–40. <https://doi.org/10.1017/s0094837300016134>.
- Frankfurter, F. (1947, May). Some reflections on the reading of statutes. *Columbia Law Review*, 47(4), 527–46.
- Gaba, J. (1999). Environmental ethics and our moral relationship to future generations: Future rights and present virtue. *Columbia Journal of Environmental Law*, 24(2), 249–88.
- Gallup, J. L., & Sachs, J. D. (2001, January–February). The economic burden of malaria. *American Journal of Tropical Medicine and Hygiene*, 64(1–2 Suppl), 85–96. <https://doi.org/10.4269/ajtmh.2001.64.85>.
- Galway-Witham, J., & Stringer, C. (2018, June 22). How did *Homo sapiens* evolve? *Science*, 360(6395), 1296–8. <https://doi.org/10.1126/science.aat6659>.
- Gilovich, T., Griffin, D., & Kahneman, D. (Eds.) (2002). *Heuristics and biases: The psychology of intuitive judgment*. Cambridge, UK: Cambridge University Press.
- Giuliano, P., & Nunn, N. (2021, July). Understanding cultural persistence and change. *The Review of Economic Studies*, 88(4), 1541–81. <https://doi.org/10.1093/restud/rdaa074>.
- Gollier, C. (2012). *Pricing the planet's future: The economics of discounting in an uncertain world*. Princeton: Princeton University Press.
- Greaves, H. (2017, November). Discounting for public policy: A survey. *Economics & Philosophy*, 33(3), 391–439. <https://doi.org/10.1017/S0266267117000062>.
- Greaves, H., & MacAskill, W. (2019, September). *The case for strong longtermism* (GPI Working Paper No. 7-2019). Global Priorities Institute. https://globalprioritiesinstitute.org/wp-content/uploads/2020/Greaves_MacAskill_strong_longtermism.pdf.
- Greaves, H., & MacAskill, W. (2021, June). *The case for strong longtermism* (GPI Working Paper No. 5-2021). Global Priorities Institute. <https://globalprioritiesinstitute.org/wp>

- content/uploads/The-Case-for-Strong-Longtermism-GPI-Working-Paper-June-2021-2.pdf.
- Greene, J., & Baron, J. (2001, June 29). Intuitions about declining marginal utility. *Journal of Behavioral Decision Making*, 14(3), 243–55. <https://doi.org/10.1002/bdm.375>.
- Greene, J. D. (2013). *Moral tribes: Emotion, reason, and the gap between us and them*. New York: Penguin Press.
- Greene, J. D. (2014). Beyond point-and-shoot morality: Why cognitive (neuro) science matters for ethics. *Ethics*, 124(4), 695–726. <https://www.journals.uchicago.edu/doi/abs/10.1086/675875>.
- Gruen, L. (2017, August 23). The moral status of animals. In E. N. Zalta (Ed.), *Stanford encyclopedia of philosophy*. <https://plato.stanford.edu/entries/moral-animal/>.
- Guthrie, C., Rachlinski, J. J., & Wistrich, A. J. (2000). Inside the judicial mind. *Cornell Law Review*, 86(4), 777–830.
- Harris v. U.S. Dep't of Veterans Affairs, 776 F.3d 907 (2015).
- Harrod, R. F. (1948). *Towards a dynamic economics: Some recent developments of economic theory and their application to policy*. London: MacMillan and Company.
- Housey Pharmaceuticals, Inc. v. Astrazeneca UK Ltd., No. 03–1193 (Fed. Cir. May 7, 2004).
- Interpretation Act (Cap 1) s 9A (1993).
- John, T., & MacAskill, W. (2020, October). *Longtermist institutional reform* (GPI Working Paper No. 14-2020). Global Priorities Institute. https://globalprioritiesinstitute.org/wp-content/uploads/Tyler-M-John-and-William-MacAskill_Longtermist-institutional-reform.pdf.
- Jowett, Inc. v. United States, 234 F.3d 1365 (2000).
- Kahneman, D., Ritov, I., Schkade, D., Sherman, S. J., & Varian, H. R. (1999). Economic preferences or attitude expressions? An analysis of dollar responses to public issues. In B. Fischhoff & B. Manski (Eds.), *Elicitation of preferences* (pp. 203–42). New York: Springer.
- Klapper, S., Schmidt, S., & Tarantola, T. (2020). Ordinary Meaning from Ordinary People [Unpublished manuscript].
- Knobe, J., Buckwalter, W., Nichols, S., Robbins, P., Sarkissian, H., & Sommers, T. (2012). Experimental philosophy. *Annual Review of Psychology*, 63(1): 81–99. <https://doi.org/10.1146/annurev-psych-120710-100350>.
- Kuecken, M., Thuilliez, J., & Valfort, M. A. (2014, December). *Does malaria control impact education? Evidence from Roll Back Malaria in Africa* (Les Notes du G-MonD, Note 12). Paris School of Economics. <https://www.parisschoolofeconomics.eu/IMG/pdf/note12-gmond-malaria-education-africa-pse-december2014.pdf>.
- Kuran, T. (2011, February–March). Legal roots of economic underdevelopment in the Middle East. *The European Financial Review*, February – March 2011, 10–11.
- Lee, T. R. & Mouritsen, S. C. (2018). Judging Ordinary Meaning. *Yale Law Journal*, 127(4), 788–879. <https://www.yalelawjournal.org/article/judging-ordinary-meaning>.
- MacAskill, W. (2019). The definition of effective altruism. In H. Greaves & T. Pummer (Eds.), *Effective altruism: Philosophical Issues* (pp. 10–28). New York: Oxford University Press. <https://doi.org/10.1093/oso/9780198841364.001.0001>.
- MacAskill, W. (2022). *What we owe the future*. New York: Basic Books.
- MacCormick, D. N., & Summers, R. S. (Eds.). (2016). *Interpreting statutes: A comparative study*. New York: Routledge.
- Martínez, E., & Tobia, K. P. (2022). *What do law professors believe about law and the legal academy? An empirical inquiry* https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4182521.

- Martínez, E., & Winter, C. K. (2021a, August 20). *Protecting future generations: A global survey of legal academics* (LPP Working Paper Series No. 1). Legal Priorities Project. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3931304.
- Martínez, E., & Winter, C. (2021b, November 26). Protecting sentient artificial intelligence: A survey of lay intuitions on standing, personhood, and general legal protection. *Frontiers in Robotics and AI: Ethics in Robotics and Artificial Intelligence*, 8, 788355. <https://doi.org/10.3389/frobt.2021.788355>.
- Martínez, E., & Winter, C. (2022). *The ordinary meaning of existential risk* [unpublished manuscript].
- Martínez, E., & Winter, C. (in press a). Cross-cultural perceptions of rights for future generations. In K. Tobia (Ed.), *Cambridge handbook of experimental jurisprudence*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4262891.
- Martínez, E., & Winter, C. (in press b). The intuitive appeal of legal longtermism. In J. Barrett, D. Thorstad, & H. Greaves (Eds.), *Essays on longtermism*. Oxford: Oxford University Press.
- Mischel, W., & Ebbesen, E. B. (1970). Attention in delay of gratification. *Journal of Personality & Social Psychology*, 16(2), 329–37. <https://doi.org/10.1037/h0029815>.
- Mogensen, A. (2019, September). *Maximal cluelessness* (GPI Working Paper No. 2-2019). Global Priorities Institute. https://globalprioritiesinstitute.org/wp-content/uploads/2020/Andreas_Mogensen_maximal_cluelessness.pdf.
- Mogensen, A. L. (2016). Contingency anxiety and the epistemology of disagreement. *Pacific Philosophical Quarterly*, 97(4), 590–611. <https://doi.org/10.1111/papq.12099>.
- Moskal v. United States, 498 U.S. 103 (1990).
- Mueller, G. O., trans. (1961), *The German Penal Code of 1871*, 107, South Hackensack, NJ: Rothman.
- Nix v. Hedden, 149 U.S. 304 (1893).
- Occidental Life Insurance Co. of Cal. v. United States, 250 F. Supp. 130 (S.D. Cal. 1965).
- O'Donoghue, T., & Rabin, M. (1999, March). Doing it now or later. *American Economic Review*, 89(1), 103–24. <https://doi.org/10.1257/aer.89.1.103>.
- O'Donoghue, T., & Rabin, M. (2015, May). Present bias: Lessons learned and to be learned. *American Economic Review*, 105(5), 273–9. <https://doi.org/10.1257/aer.p20151085>.
- Ord, T. (2020). *The precipice: Existential risk and the future of humanity*. New York: Hachette Books.
- Parfit, D. (1984). *Reasons and persons*. Oxford: Oxford University Press.
- Pettit, P., & Smith, M. (2000, June). Global consequentialism. In B. Hooker, E. Mason, & D. E. Miller (Eds.), *Morality, rules, and consequences: A critical reader* (pp. 121–33). Edinburgh: Edinburgh University Press.
- Pigou, A. C. (1932, December 1). The effect of reparations on the ratio of international interchange. *The Economic Journal*, 42(168): 532–43. <https://doi.org/10.2307/2223778>.
- Pirie, F. (2021). *The rule of laws: A 4,000-year quest to order the world*. New York: Basic Books.
- Pogge, T. (1989). *Realizing Rawls*. Ithaca: Cornell University Press.
- Pryce, J., Richardson, M., & Lengeler, C. (2018). Insecticide-treated nets for preventing malaria. *Cochrane Database of Systematic Reviews*, 11. <https://doi.org/10.1002/14651858.CD000363.pub3>.
- Ramsey, F. P. (1928, December). A mathematical theory of saving. *The Economic Journal*, 38(152), 543–59. <https://doi.org/10.2307/2224098>.
- Rawls, J. (1971). *A theory of justice*. Cambridge, MA: Harvard University Press.
- Richards v. United States, 369 U.S. 1 (1962).
- River Wear Commissioners v Adamson, 2 App Cas 743 (1877).

- Roser, M., Ritchie, H., & Ortiz-Ospina, E. (2019, May). *World population growth*. Our World in Data. <https://ourworldindata.org/world-population-growth>.
- Sachs, J., & Malaney, P. (2002). The economic and social burden of malaria. *Nature*, 415(6872), 680–5. <https://doi.org/10.1038/415680a>.
- Schlebusch, C. M., Malmström, H., Günther, T., Sjödin, P., Coutinho, A., Edlund, H., Jakobsson, M. (2017, November 3). Southern African ancient genomes estimate modern human divergence to 350,000 to 260,000 years ago. *Science*, 358(6363), 652–5. <https://doi.org/10.1126/science.aao6266>.
- Schubert, S., Caviola, L., & Faber, N. S. (2019). The psychology of existential risk: Moral judgments about human extinction. *Scientific Reports*, 9(1), 1–8. <https://doi.org/10.1038/s41598-019-50145-9>.
- Schwarz, N., Bless, H., Strack, F., Klumpp, G., Rittenauer-Schatka, H., & Simons, A. (1991). Ease of retrieval as information: Another look at the availability heuristic. *Journal of Personality & Social Psychology*, 61(2), 195–202. <https://doi.org/10.1037/0022-3514.61.2.195>.
- Sidgwick, H. (1907), *The methods of ethics* (7th ed.). London: Macmillan.
- Singer, P. (1972). *Famine, affluence, and morality*. Oxford: Oxford University Press.
- Singer, P. (1973). Animal liberation. In R. Garner (Ed.), *Animal rights* (pp. 7–18). London: Palgrave Macmillan. https://doi.org/10.1007/978-1-349-25176-6_1.
- Singer, P. ([1981] 2011). *The expanding circle: Ethics, evolution, and moral progress*. Princeton & Oxford: Princeton University Press.
- Slocum, B. G. (2015). *Ordinary meaning: A theory of the most fundamental principle of legal interpretation*. Chicago: University of Chicago Press. <https://doi.org/10.7208/chicago/9780226304991.001.0001>.
- Slocum, B. G., & Wong, J. (2021). The Vienna Convention and the ordinary meaning of international law. *Yale Journal of International Law*, 46(2), 192–239. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3556892.
- Slovic, P. (2010). The more who die, the less we care. In P. Slovic (Ed.), *The feeling of risk: New perspectives on risk perception* (pp. 97–106). New York: Routledge.
- Solow, R. M. (1974). The economics of resources or the resources of economics. In C. Gopalakrishnan (Ed.), *Classic papers in natural resource economics* (pp. 257–76). London: Palgrave MacMillan. https://doi.org/10.1057/9780230523210_13.
- Stern, N. (2008). The economics of climate change. *American Economic Review*, 98(2), 1–37. <https://doi.org/10.1257/aer.98.2.1>.
- Tarsney, C. (2020, May). *The epistemic challenge to longtermism* (GPI Working Paper No. 10-2019). Global Priorities Institute. <https://globalprioritiesinstitute.org/wp-content/uploads/Working-paper-10-Christian-Tarsney.pdf>.
- Tetlock, P. E., & Gardner, D. (2016), *Superforecasting: The art and science of prediction*. New York: Crown Publishers.
- Tobia, K. P. (2020a, February 10). *Legal concepts and legal expertise* [Unpublished manuscript]. https://privpapers.ssrn.com/sol3/papers.cfm?abstract_id=3536564.
- Tobia, K. P. (2020b). Testing ordinary meaning. *Harvard Law Review*, 134(2), 726–806. <https://harvardlawreview.org/2020/12/testing-ordinary-meaning/>.
- Tobia, K. P. (2022). Experimental jurisprudence. *University of Chicago Law Review*, 89(3), 735–802. <https://lawreview.uchicago.edu/publication/experimental-jurisprudence>.
- Todd, B. (2017, October). Why your impact in millions of years could be what most matters. 80,000 Hours. <https://80000hours.org/articles/future-generations/>.
- Trivers, R. L. (1971). The evolution of reciprocal altruism. *The Quarterly Review of Biology*, 46(1), 35–57. <https://www.journals.uchicago.edu/doi/10.1086/406755>.

- Tversky, A., & Kahneman, D. (1973, September). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 5(2): 207–232. [https://doi.org/10.1016/0010-0285\(73\)90033-9](https://doi.org/10.1016/0010-0285(73)90033-9).
- United Nations (2019). 2019 Revision of World Population Prospects. <https://population.un.org/wpp/>.
- United States v. Moore, 572 F.3d 489 (2009).
- United States v. Shafer, 573 F.3d 267 (2009).
- United States v. Turkette, 452 U.S. 576 (1981).
- Venter v. R, TS 910 (1907).
- Vienna Convention on the Law of Treaties art. 31, opened for signature May 23, 1969, 1155 U.N.T.S. 331.
- Watson, A. (1991), *Roman Law and comparative law*. Athens: University of Georgia Press.
- Waytz, A., Iyer, R., Young, L., & Graham, J. (2016, February 17). Ideological differences in the expanse of empathy. In P. Valdesolo & J. Graham (Eds.), *Social psychology of political polarization* (pp. 61–77). New York: Routledge.
- Waytz, A., Iyer, R., Young, L., Haidt, J., & Graham, J. (2019, September 26). Ideological differences in the expanse of the moral circle. *Nature Communications*, 10(1), 4389. <https://www.nature.com/articles/s41467-019-12227-0>.
- Wedgwood, R. (2007, October). Normativism defended. In B. P. McLaughlin & J. Cohen (Eds.), *Contemporary debates in philosophy of mind* (pp. 85–101). Blackwell Publishing.
- Winter, C. K. (2020, February 10). The value of behavioral economics for EU judicial decision-making. *German Law Journal*, 21(2), 240–64. <https://doi.org/10.1017/glj.2020.3>.
- Winter, C. K. (in press). The challenges of artificial judicial decision-making for liberal democracy. In Bystranowski, P., Janik, B., & Próchnicki, M. (Eds.), *Judicial decision-making: Integrating empirical and theoretical perspectives*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3933648.
- Winter, C. K., Schuett, J., Martínez, E., Van Arsdale, S., Araújo, R., Hollman, N., Rotola, G. (2021, January). Legal priorities research: A research agenda. Legal Priorities Project. https://www.legalpriorities.org/research_agenda.pdf.
- Yudkowsky, E. (2008). Cognitive biases potentially affecting judgment of global risks. In *Global catastrophic risks* (pp. 91–119). Oxford: Oxford University Press. <https://doi.org/10.1093/oso/9780198570509.001.0001>.